

Claims (for amendment filed February 19, 2005)

44. (previously presented) A processor device comprising

an instruction stream transformation unit that transforms code blocks of instructions from an original instruction set architecture to a transformed instruction set architecture,

a regular cache that stores instructions in said original instruction set architecture,

an instruction stream cache that stores instructions in said transformed instruction set architecture, and

an execute unit for executing instructions,

wherein said instruction stream transformation unit transforms said code blocks of instructions from said original instruction set architecture to said transformed instruction set architecture,

wherein said instruction stream cache stores said code blocks after transformation to said transformed instruction set architecture for possible future execution,

wherein said instruction stream cache comprises means of storing and later fetching a transformed code block that spans more than one cache line in said instruction stream cache,

wherein said instruction stream cache is addressed by some of said fetch requests from said execute unit and can potentially respond to some of said fetch requests for said code blocks after transformation without requiring cache hit information from said regular cache after said code blocks have already been transformed and stored into the instruction stream cache, and

whereby the execution of a program code by said processor device is accelerated by transforming portions of said program code at run-time into the transformed instruction set architecture for more efficient execution and caching the transformed code within the instruction stream cache for possible repeated execution without requiring repeated transformations.

45. (previously presented) A processor device as in claim 44 wherein

said execute unit can directly execute instructions in both said original instruction set architecture and said transformed instruction set architecture,

Methods for Increasing Instruction-Level Parallelism

whereby the execute unit can execute code in said original instruction set architecture without having to first wait for said code to be transformed into said transformed instruction set architecture.

46. (previously presented) A processor device as in claim 45 wherein said execute unit is an in-order execute unit that retires instructions and commits their results in the same order as the instructions occur in the code.

47. (previously presented) A processor device as in claim 45 wherein said execute unit is a dynamic out-of-order execute unit that can retire instructions out of order compared to the order of the instructions in the code.

48. (previously presented) A processor device as in claim 44 further comprising a working memory connected to said instruction stream transformation unit for storing intermediate calculations during the process of transforming code from said original instruction set architecture into said transformed instruction set architecture.

49. (previously presented) A processor device as in claim 44 wherein said instruction stream transformation unit transforms blocks of code which are presumed hyper-blocks.

50. (previously presented) A processor device as in claim 44 wherein said instruction stream transformation unit transforms blocks of code which are denoted as hyper-blocks in original instruction set architecture code.

51. (previously presented) A processor device as in claim 44 wherein said instruction stream cache comprises means of using a tag for each cache line to denote whether the tag is a start of a hyper-block.

52. (previously presented) A processor device as in claim 44 wherein said instruction stream cache comprises means of using a hyper-block ID as an alternative way of addressing a transformed block of code.

Methods for Increasing Instruction-Level Parallelism

53. (previously presented) A processor device as in claim 52 wherein said instruction stream cache comprises means of storing a plurality of hyper-block lines that are chained together using a common hyper-block ID plus pointers.

54. (previously presented) A processor device as in claim 44 wherein said instruction stream cache only hits when a transformed block of code is entered starting from the first instruction of the block.

55. (previously presented) A processor device as in claim 44 wherein said instruction stream transformation unit comprises means of instruction re-ordering.

56. (previously presented) A processor device as in claim 44 wherein said instruction stream transformation unit comprises means of performing predication if-conversion to convert an if-then-else construct in said code into a predication calculation and a series of predicated instructions that conditionally commit their results depending on the results of said predication calculation.

57. (previously presented) A processor device as in claim 44 wherein said instruction stream transformation unit comprises means of converting a load instruction into a speculative load and a load activation instruction pair,

whereby the possibility of more efficient scheduling of the transformed code is enabled by allowing the speculative load to be scheduled earlier than a normal load could be scheduled without this conversion thus minimizing possible waiting during execution for this memory load to complete.

58. (previously presented) A processor device as in claim 44 wherein said instruction stream transformation unit comprises a parallel dependency detector circuit for detecting possible dependencies between instructions,

Methods for Increasing Instruction-Level Parallelism

whereby said instruction stream transformation unit can efficiently detect said possible dependencies during the process of instruction stream transformation.

59. (previously presented) A processor device as in claim 44 wherein said instruction stream transformation unit transforms said code blocks by dividing long instruction sequences into sequences of a defined maximum number of instructions called an instruction window that is equal to a number of instructions that said instruction stream transformation unit can work with effectively.

60. (previously presented) A processor device as in claim 59 wherein said instruction stream transformation unit transforms said code blocks by dividing long instruction sequences into overlapping sequences of a defined maximum number of instructions called overlapping instruction windows,

whereby this enables the efficient scheduling of code both within the middle of instruction windows and within the overlap regions between instruction windows.

61. (previously presented) A processor device as in claim 44 wherein said instruction stream transformation unit comprises means of creating a dependency matrix to represent potential dependencies between instructions,

whereby the dependency matrix provides dependency information in an efficiently accessed manner for the instruction stream transformation unit.

62. (previously presented) A processor device as in claim 61 wherein said instruction stream transformation unit comprises means of creating an operand mapping table to represent potential writes of operands by instructions,

whereby said operand mapping table enables dependencies between instructions to be detected efficiently and for the dependency matrix to be built efficiently by said instruction stream transformation unit.

Methods for Increasing Instruction-Level Parallelism

63. (previously presented) A processor device as in claim 62 wherein said instruction stream transformation unit comprises means of performing register renaming to rename uses of registers in the transformed code blocks when necessary to minimize write-after-write hazards between write instructions,

whereby scheduling dependencies caused by said write-after-write hazards is reduced.

64. (previously presented) A processor device as in claim 63 wherein said register renaming by the instruction stream transformation unit allocates a set of physical registers that are separate from a second set of registers allocated by dynamic register renaming done by the execute unit,

whereby said instruction stream transformation unit and said execute unit are each able to independently allocate physical registers for register renaming without conflicting with each other.

65. (previously presented) A processor device as in claim 44 wherein the instruction stream transformation unit comprises means of performing an instruction scheduling algorithm to schedule a code block.

66. (previously presented) A processor device as in claim 65 wherein said scheduling algorithm is a list scheduling algorithm comprising the steps of

doing a basic forward iterative traversal to calculate the minimum cycle number of each instruction in said block from the start of said block,

propagating the depth of each leaf instruction to all its predecessors to calculate each instruction's priority as defined by the depth of its deepest child, and

performing a second forward iterative traversal to schedule instructions for execution cycle-by-cycle, where the scheduling priorities are used in conjunction with a ready list of instructions that are ready to be scheduled because all dependencies have been resolved.

Methods for Increasing Instruction-Level Parallelism

67. (previously presented) A processor device as in claim 44 wherein

said execute unit comprises means of performing dynamic memory disambiguation,

said transformed instruction set architecture supports notations for indicating ambiguous memory operations,

and said instruction stream transformation unit calculates said notations for indicating ambiguous memory operations.

68. (previously presented) A processor device as in claim 67 wherein said instruction stream transformation unit comprises means of creating an ambiguous memory dependency matrix.

69. (previously presented) A processor device as in claim 67 wherein said instruction stream transformation unit comprises means of converting an ambiguous memory read instruction into a speculative read and a read check instruction pair.

70. (previously presented) A processor device as in claim 44 comprising a dynamically-scheduled execute unit.

71. (previously presented) A processor device as in claim 44

wherein the transformed instruction set architecture comprises dependency notation means of explicitly describing dependencies between instructions, and

wherein the instruction stream transformation unit calculates dependency notations for said code blocks during the process of transforming said code blocks into said transformed instruction set architecture,

whereby the transformed code blocks can be executed without having to redetect all static dependencies.

Methods for Increasing Instruction-Level Parallelism

72. (previously presented) A processor device as in claim 71 wherein said means of explicitly describing dependency notation means allows instructions to be grouped into mini-tuples for dependency notations.

73. (previously presented) A processor device as in claim 71 wherein said dependencies are represented by dependency pointers in the transformed instruction set architecture.

74. (previously presented) A processor device as in claim 71 wherein said dependencies are represented by dependency vectors in the transformed instruction set architecture.

75. (previously presented) A processor device as in claim 44 further comprising means to perform semi-dynamic instruction code re-writing and re-scheduling,

whereby the code can be further optimized based on run-time information compared to code that is transformed only once through the instruction stream transformation unit.

76. (previously presented) A processor device as in claim 44 further comprising at least one run-time history table.

77. (previously presented) A processor device as in claim 44 comprising a value prediction history table to record the history of success or failure of previous executions of value predictions,

whereby run-time behavior about value predictions can be recorded in order to help optimize subsequent instruction scheduling.

78. (currently amended) A processor device as in claim 44 comprising a predicate history table to record the history of previous executions of predicate calculation instructions,

whereby run-time behavior about predicates can be recorded in order to help optimize subsequent instruction scheduling.

Methods for Increasing Instruction-Level Parallelism

79. (previously presented) A processor device as in claim 44 comprising

a data cache and

a data hit/miss history table to record the history of whether previous executions of memory access instructions were hits or misses in said data cache,

whereby run-time behavior about data hit/misses can be recorded in order to help optimize subsequent instruction scheduling.

80. (previously presented) A processor device comprising an ambiguous memory conflict history table to record the history of whether previous executions of promoted ambiguous read instructions cause memory conflicts or not,

wherein said processor device uses the history recorded in said ambiguous memory conflict history table as a predictive indicator of the likelihood of memory conflicts in subsequent executions in order to affect the scheduling order of said promoted ambiguous read instructions relative to other instructions in subsequent executions,

whereby run-time behavior about ambiguous memory conflicts can be recorded in order to help optimize subsequent instruction scheduling.

81. (previously presented) A method of providing precise interrupts in a processor implementing instruction stream transformation comprising the steps of

mapping an original instruction set architecture instruction to an equivalent group of one or more transformed instruction set architecture instruction(s),

using physical registers that have not been committed to logical registers to hold results, and

allowing the final instruction in said group to commit the physical register result(s) to logical register(s).

Methods for Increasing Instruction-Level Parallelism

82. (previously presented) A method of providing precise interrupts in a processor implementing instruction stream transformation comprising the steps of

assigning an instruction sequence number to each original instruction set architecture instruction starting from the beginning of the code block,

marking each transformed instruction set architecture instruction with the corresponding instruction sequence number, and

committing the results of instructions in order of the instruction sequence numbers.

83. (cancelled)

84. (currently amended) A processor device as in claim 44 comprising

means of executing instructions from an instruction set architecture

wherein the instruction set architecture ~~can~~ explicitly ~~note~~ notes dependencies between instructions by using dependency vectors,

and wherein said processor device stores said instructions including said dependency vectors for possible repeated execution.

85. (currently amended) A processor device as in claim 44 comprising

means of executing instructions from an instruction set architecture

wherein the instruction set architecture ~~can~~ explicitly ~~note~~ notes dependencies between instructions by using dependency pointers,

and wherein said processor device stores said instructions including said dependency pointers for possible repeated execution.

86. (cancelled)

Methods for Increasing Instruction-Level Parallelism

87. (previously presented) A processor device as in claim 44 wherein said means of storing and later fetching said transformed code block that spans more than one cache line in said instruction stream cache comprises means of using data structures to associate together multiple cache lines that each store a portion of said transformed code block.